


from library



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 0 981 113 A2**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
23.02.2000 Bulletin 2000/08

(51) Int. Cl.⁷: **G07D 7/00**

(21) Application number: 99113163.2

(22) Date of filing: 07.07.1999

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE**
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Rhoads, Geoffrey B.
West Linn, Oregon 97068 (US)

(74) Representative:
Meddle, Alan Leonard et al
FORRESTER & BOEHMERT
Franz-Joseph-Strasse 38
80801 München (DE)

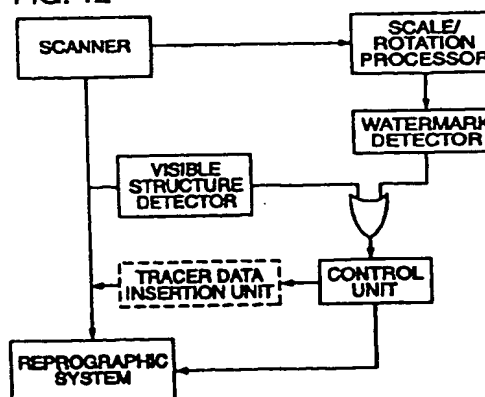
(30) Priority: 31.07.1998 US 127502

(71) Applicant: Digimarc Corporation
Lake Oswego, OR 97035 (US)

(54) **Digital watermarking and methods for security documents**

(57) Security documents (e.g. passports, currency, event tickets, and the like) are encoded to convey machine-readable multi-bit binary information (e.g. digital watermark), usually in a manner not alerting human viewers that such information is present. The documents can be provided with overt or subliminal calibration patterns. When a document incorporating such a pattern is scanned (e.g. by a photocopier), the pattern facilitates detection of the encoded information notwithstanding possible sealing or rotation of the scan data. The calibration pattern can serve as a carrier for the watermark information, or the watermark can be encoded independently. In one embodiment, the watermark and the calibration pattern are formed on the document by an intaglio process, with or without ink. A photocopier responsive to such markings can take predetermined action if reproduction of a security document is attempted. A passport processing station responsive to such markings can use the decoded binary data to access a database having information concerning the passport holder. Some such apparatuses detect both the watermark data and the presence of a visible structure characteristic of a security document (e.g., a printed seal of the document's issuer).

FIG. 12



EP 0 981 113 A2

Description

Related Application Data

[0001] This application is related to commonly-owned patents 5,862,260, 5,850,481, 5,841,978, 5,841,886, 5,832,119, 5,822,436, 5,809,160, 5,745,604, 5,636,292, 5,710,834, 5,721,788, 5,748,763, 5,748,783, and 5,768,426.

Field of the Invention

[0002] The present invention relates to methods and systems for inconspicuously embedding binary data in security documents, and associated methods/systems for detecting/decoding such data. ("Security document" is used herein to refer to negotiable financial instruments (e.g. banknotes, travelers checks, bearer bonds), passports, visas, other immigration documents, stock certificates, postal stamps, lottery tickets, sports/concert tickets, etc.) One application of this the invention is in discouraging counterfeiting of security documents. Another is in transferring machine-readable information through such documents, without alerting human viewers to the presence of such information.

Background and Summary of the Invention

[0003] Digital watermarking (sometimes termed "data hiding" or "data embedding") is a growing field of endeavor, with several different approaches. The present assignee's work is reflected in the patents and applications detailed above, together with laid-open PCT application WO97/43736. Other work is illustrated by U.S. patents 5,734,752, 5,646,997, 5,659,726, 5,664,018, 5,671,277, 5,687,191, 5,687,236, 5,689,587, 5,568,570, 5,572,247, 5,574,962, 5,579,124, 5,581,500, 5,613,004, 5,629,770, 5,461,426, 5,743,631, 5,488,664, 5,530,759, 5,539,735, 4,943,973, 5,337,361, 5,404,160, 5,404,377, 5,315,098, 5,319,735, 5,337,362, 4,972,471, 5,161,210, 5,243,423, 5,091,966, 5,113,437, 4,939,515, 5,374,976, 4,855,827, 4,876,617, 4,939,515, 4,963,998, 4,969,041, and published foreign applications WO 98/02864, EP 822,550, WO 97/39410, WO 96/36163, GB 2,196,167, EP 777,197, EP 736,860, EP 705,025, EP 766,468, EP 782,322, WO 95/20291, WO 96/26494, WO 96/36935, WO 96/42151, WO 97/22206, WO 97/26733. Some of the foregoing patents relate to visible watermarking techniques. Other visible watermarking techniques (e.g. data glyphs) are described in U.S. patents 5,706,364, 5,689,620, 5,684,885, 5,680,223, 5,668,636, 5,640,647, 5,594,809.

[0004] Much of the work in data embedding is not in the patent literature but rather is published in technical articles. In addition to the patentees of the foregoing patents, some of the other workers in this field (whose

watermark-related writings can be found by an author search in the INSPEC or NEXIS databases, among others) include I. Pitas, Eckhard Koch, Jian Zhao, Norishige Morimoto, Laurence Boney, Kineo Matsui, A.Z. Tirkel, Fred Mintzer, B. Macq, Ahmed H. Tewfik, Frederic Jordan, Naohisa Komatsu, Joseph O'Ruanaidh, Neil Johnson, Ingemar Cox, Minerva Yeung, and Lawrence O'Gorman.

[0005] The artisan is assumed to be familiar with the foregoing prior art.

[0006] In the following disclosure it should be understood that references to watermarking encompass not only the assignee's watermarking technology, but can likewise be practiced with any other watermarking technology, such as those indicated above.

[0007] Watermarking can be applied to myriad forms of information. The present disclosure focuses on its applications to security documents. However, it should be recognized that the principles discussed below can also be applied outside this area.

[0008] Most of the prior art in image watermarking has focused on pixelated imagery (e.g. bit-mapped images, JPEG/MPEG imagery, VGA/SVGA display devices, etc.). In most watermarking techniques, the luminance or color values of component pixels are slightly changed to effect subliminal encoding of binary data through the image. (This encoding can be done directly in the pixel domain, or after the signal has been processed and represented differently - e.g. as DCT or wavelet coefficients, or as compressed data, etc.)

[0009] While pixelated imagery is a relatively recent development, security documents—commonly employing line art—go back centuries. One familiar example is U.S. paper currency. On the one dollar banknote, for example, line art is used in several different ways. One is to form intricate webbing patterns (sometimes termed "guilloche patterns") around the margin of the note (generally comprised of light lines on dark background). Another is to form gray scale imagery, such as the portrait of George Washington (generally comprised of dark lines on a light background).

[0010] There are two basic ways to simulate greyscales in security document line art. One is to change the relative spacings of the lines to effect a lightening or darkening of an image region. Fig. 1A shows such an arrangement; area B looks darker than area A due to the closer spacings of the component lines. The other technique is to change the widths of the component lines—wider lines resulting in darker areas and narrower lines resulting in lighter areas. Fig. 1B shows such an arrangement. Again, area B looks darker than area A, this time due to the greater widths of the component lines. These techniques are often used together. Ultimately, a given region simply has more or less ink.

[0011] In my application 08/438,159 I introduced, and in my application 09/074,034 I elaborated on, techniques for watermarking line art by making slight changes to the widths, or positions, of the component

lines. Such techniques are further expanded in the present disclosure.

[0012] In several of my cited patents, I discussed various "calibration signals" that can be used to facilitate the decoding of watermark data despite corruption of the encoded image, such as by scaling or rotation. (See, e.g., patent 5,832,119.) Common counterfeiting techniques — e.g. color photocopying, or scanning/jet printing — often introduce such corruption, whether deliberately or accidentally. Accordingly, it is important that watermarks embedded in security documents be detectable notwithstanding such effects. Calibration signals particularly suited for use with security documents are detailed in this disclosure.

[0013] In accordance with embodiments of the present invention, security documents are encoded to convey machine-readable multi-bit binary information (e.g. digital watermarks), usually in a manner not alerting human viewers that such information is present. The documents can be provided with overt or subliminal calibration patterns. When a document incorporating such a pattern is scanned (e.g. by a photocopier), the pattern facilitates detection of the encoded information notwithstanding possible scaling or rotation of the scan data. The calibration pattern can serve as a carrier for the watermark information, or the watermark can be encoded independently. In one embodiment, the watermark and the calibration pattern are formed on the document by an intaglio process, with or without ink. A photocopier responsive to such markings can take predetermined action if reproduction of a security document is attempted. A passport processing station responsive to such markings can use the decoded binary data to access a database having information concerning the passport holder. Some such apparatuses detect both the watermark data and the presence of a visible structure characteristic of a security document.

[0014] The foregoing and other features and advantages of the present technology will be more readily apparent from the following detailed description, which proceeds with reference to the accompanying drawings.

Brief Description of the Drawings

[0015]

Figs. 1A and 1B show prior art techniques for achieving grey-scale effects using line art.

Fig. 2 shows a virtual array of grid points that can be imposed on a security document image according to an embodiment of the present invention.

Fig. 3 shows a virtual array of regions that can be imposed on a security document image according to the Fig. 2 embodiment.

Fig. 4 shows an excerpt of Fig. 3 with a line from a line art image passing therethrough.

Fig. 5 shows changes to the width of the line of Fig.

3 to effect watermark encoding.

Fig. 6 shows changes to the position of the line of Fig. 3 to effect watermark encoding.

Figs. 7A and 7B show aspects of watermark and calibration blocks according to an embodiment of the invention.

Fig. 8 shows an illustrative reference grey-scale calibration tile.

Figs. 9A-9C show steps in the design of a weave calibration pattern according to an embodiment of the invention.

Fig. 10 shows the generation of error data used in designing a weave calibration pattern according to an embodiment of the invention.

Fig. 11 is a block diagram of a passport processing station according to another embodiment of the invention.

Fig. 12 is a block diagram of a photocopier according to another embodiment of the invention.

Detailed Description

[0016] By way of introduction, the present specification begins with review of techniques for embedding watermark data in line art, as disclosed in my application 09/074,034.

[0017] Referring to Fig. 2, the earlier-described technique employs a grid 10 of imaginary reference points arrayed over a line art image. The spacing between points is 250 microns in the illustrated arrangement, but greater or lesser spacings can of course be used.

[0018] Associated with each grid point is a surrounding region 12, shown in Fig. 3. As described below, the luminosity (or reflectance) of each of these regions 12 is slightly changed to effect subliminal encoding of binary data.

[0019] Region 12 can take various shapes; the illustrated rounded-rectangular shape is representative only. (The illustrated shape has the advantage of encompassing a fairly large area while introducing fewer visual artifacts than, e.g., square regions.) In other embodiments, squares, rectangles, circles, ellipses, etc., can alternatively be employed.

[0020] Fig. 4 is a magnified view of an excerpt of Fig. 3, showing a line 14 passing through the grid of points. The width of the line, of course, depends on the particular image of which it is a part. The illustrated line is about 40 microns in width; greater or lesser widths can naturally be used.

[0021] In one encoding technique, shown in Fig. 5, the width of the line is controllably varied so as to change the luminosity of the regions through which it passes. To increase the luminosity (or reflectance), the line is made narrower (i.e. less ink in the region). To decrease the luminosity, the line is made wider (i.e. more ink).

[0022] Whether the luminance in a given region should be increased or decreased depends on the particular watermarking algorithm used. Any algorithm can

be used, by changing the luminosity of regions 12 as the algorithm would otherwise change the luminance or colors of pixels in a pixelated image. (Some watermarking algorithms effect their changes in a transformed domain, such as DCT, wavelet, or Fourier. However, such changes are ultimately manifested as changes in luminance or color.)

[0023] In an exemplary algorithm, the binary data is represented as a sequence of -1s and 1s, instead of 0s and 1s. (The binary data can comprise a single datum, but more typically comprises several. In an illustrative embodiment, the data comprises 128 bits, some of which are error-correcting or -detecting bits.)

[0024] Each element of the binary data sequence is then multiplied by a corresponding element of a pseudo-random number sequence, comprised of -1s and 1s, to yield an intermediate data signal. Each element of this intermediate data signal is mapped to a corresponding sub-part of the image, such as a region 12. (Commonly, each element is mapped to several such sub-parts.) The image in (and optionally around) this region is analyzed to determine its relative capability to conceal embedded data, and a corresponding scale factor is produced. Exemplary scale factors may range from 0 to 3. The scale factor for the region is then multiplied by the element of the intermediate data signal mapped to the region in order to yield a "tweak" or "bias" value for the region. In the illustrated case, the resulting tweaks can range from -3 to 3. The luminosity of the region is then adjusted in accordance with the tweak value. A tweak value of -3 may correspond to a -5% change in luminosity; -2 may correspond to -2% change; -1 may correspond to -1% change; 0 may correspond to no change; 1 may correspond to +1% change; 2 may correspond to +2% change, and 3 may correspond to +5% change. (This example follows the basic techniques described in the Real Time Encoder embodiment disclosed in patent 5,710,834.)

[0025] In Fig. 5, the watermarking algorithm determined that the luminance of region A should be reduced by a certain percentage, while the luminance of regions C and D should be increased by certain percentages.

[0026] In region A, the luminance is reduced by increasing the line width. In region D, the luminance is increased by reducing the line width; similarly in region C (but to a lesser extent).

[0027] No line passes through region B, so there is no opportunity to change the region's luminance. This is not fatal to the method, however, since the exemplary watermarking algorithm redundantly encodes each bit of data in sub-pans spaced throughout the line art image.

[0028] The changes to line widths in regions A and D of Fig. 5 are exaggerated for purposes of illustration. While the illustrated variance is possible, most implementations will typically modulate the line width 3 - 50% (increase or decrease).

[0029] (Many watermarking algorithms routinely oper-

ate within a signal margin of about +/- 1% changes in luminosity to effect encoding. That is, the "noise" added by the encoding amounts to just 1% or so of the underlying signal. Lines typically don't occupy the full area of a region, so a 10% change to line width may only effect a 1% change to region luminosity, etc. Security documents are different from photographs in that the artwork generally need not convey photorealism. Thus, security documents can be encoded with higher energy than is used in watermarking photographs, provided the result is still aesthetically satisfactory. To illustrate, localized luminance changes on the order of 10% are possible in security documents, while such a level of watermark energy in photographs would generally be considered unacceptable. In some contexts, localized luminance changes of 20, 30, 50 or even 100% are acceptable.)

[0030] In the illustrated technique, the change to line width is a function solely of the watermark tweak (or watermark/calibration pattern tweak, as discussed below) to be applied to a single region. Thus, if a line passes through any part of a region to which a tweak of 2% is to be applied, the line width in that region is changed to effect the 2% luminance difference. In variant techniques, the change in line width is a function of the line's position in the region. In particular, the change in line width is a function of the distance between the region's center grid point and the line's closest approach to that point. If the line passes through the grid point, the full 2% change is effected. At successively greater distances, successively smaller changes are applied. The manner in which the magnitude of the tweak changes as a function of line position within the region can be determined by applying one of various interpolation algorithms, such as the bi-linear, bi-cubic, cubic splines, custom curve, etc.

[0031] In other variant techniques, the change in line width in a given region is a weighted function of the tweaks for adjoining or surrounding regions. Thus, the line width in one region may be increased or decreased in accordance with a tweak value corresponding to one or more adjoining regions.

[0032] Combinations of the foregoing techniques can also be employed.

[0033] In the foregoing techniques, it is sometimes necessary to trade-off the tweak values of adjoining regions. For example, a line may pass along a border between regions, or pass through the point equidistant from four grid points ("equidistant zones"). In such cases, the line may be subject to conflicting tweak values - one region may want to increase the line width, while another may want to decrease the line width. (Or both may want to increase the line width, but differing amounts.) Similarly in cases where the line does not pass through an equidistant zone, but the change in line width is a function of a neighborhood of regions whose tweaks are of different values. Again, known interpolation functions can be employed to determine the weight to be given the tweak from each region in determining

what change is to be made to the line width in any given region.

[0034] In the exemplary watermarking algorithm, the average change in luminosity across the security document image is zero, so no generalized lightening or darkening of the image is apparent. The localized changes in luminosity are so minute in magnitude, and localized in position, that they are essentially invisible (e.g. inconspicuous/subliminal) to human viewers.

[0035] An alternative technique is shown in Fig. 6, in which line position is changed rather than line width.

[0036] In Fig. 6 the original position of the line is shown in dashed form, and the changed position of the line is shown in solid form. To decrease a region's luminosity, the line is moved slightly closer to the center of the grid point; to increase a region's luminosity, the line is moved slightly away. Thus, in region A, the line is moved towards the center grid point, while in region D it is moved away.

[0037] It will be noted that the line on the left edge of region A does not return to its nominal (dashed) position as it exits the region. This is because the region to the left of region A also is to have decreased luminosity. Where possible, it is generally preferable not to return a line to its nominal position, but instead to permit shifted lines to remain shifted as they enter adjoining regions. So doing permits a greater net line movement within a region, increasing the embedded signal level.

[0038] Again, the line shifts in Fig. 6 are somewhat exaggerated. More typical line shifts are on the order of 3 - 50 microns.

[0039] One way to think of the Fig. 6 technique is to employ a magnetism analogy. The grid point in the center of each region can be thought of as a magnet. It either attracts or repels lines. A tweak value of -3, for example, may correspond to a strong-valued attraction force; a tweak value of +2 may correspond to a middle-valued repulsion force, etc. In Fig. 6, the grid point in region A exhibits an attraction force (i.e. a negative tweak value), and the grid point in region D exhibits a repulsion force (e.g. a positive tweak value).

[0040] The magnetic analogy is useful because the magnetic effect exerted on a line depends on the distance between the line and the grid point. Thus, a line passing near a grid point is shifted more in position than a line near the periphery of the region.

[0041] (Actually, the magnetism analogy can serve as more than a conceptual tool. Instead, magnetic effects can be modeled in a computer program and serve to synthesize a desired placement of the lines relative to the grid points. Arbitrarily customized magnetic fields can be used.)

[0042] Each of the variants applicable to Fig. 5 is likewise applicable to Fig. 6.

[0043] Combinations of the embodiments of Figs. 5 and 6 can of course be used, resulting in increased watermark energy, better signal-to-noise ratio and, in many cases, less noticeable changes.

[0044] In still a further technique, the luminance in each region is changed while leaving the line unchanged. This can be effected by sprinkling tiny dots of ink in the otherwise-vacant parts of the region. In high quality printing, of the type used with security documents, droplets on the order of 3 microns in diameter can be deposited. (Still larger droplets are still beyond the perception threshold for most viewers.) Speckling a region with such droplets (either in a regular array, or random, or according to a desired profile such as Gaussian), can readily effect a 1% or so change in luminosity. (Usually dark droplets are added to a region, effecting a decrease in luminosity. Increases in luminosity can be effected by speckling with a light colored ink, or by forming light voids in line art otherwise present in a region.) (Actually, production realities often mean that many such microdots will not print, but statistically some will.)

[0045] In a variant of the speckling technique, very thin mesh lines can be inserted in the artwork - again to slightly change the luminance of one or more regions (so-called "background tinting").

[0046] The following portion of the specification reviews a calibration, or synchronization pattern used in an illustrative security document to facilitate proper registration of the watermark data for decoding. It may be helpful to being by reviewing further details about the illustrative watermarking method.

[0047] Referring to Fig. 7A, an exemplary watermark is divided into "cells" that are 250 microns on a side, each conveying a single bit of information. The cells are grouped into a "block" having 128 cells on a side (i.e. 16,384 cells per block). The blocks are tiled across the region being watermarked (e.g. across the face of a security document).

[0048] As noted, the watermark payload consists of 128 bits of data. Each bit is represented by 128 different cells within each block. (The mapping of bits to cells can be pseudo-random, sequential, or otherwise.) The 128 "0"s and "1"s of the watermark data are randomized into substantially equal-probability "1"s and "-1"s by a pseudo-random function to reduce watermark visibility. Where a cell has a value of "1," the luminance of the corresponding area of the image is slightly increased; where a cell has a value of "-1," the luminance of the corresponding area of the image is slightly decreased (or vice versa). In some embodiments, the localized changes to image luminance due to the +1/-1 watermark cell values are scaled in accordance with data-hiding attributes of the local area (e.g. to a range of +/- 4 digital numbers) to increase the robustness of the watermark without compromising its imperceptibility.

[0049] It should be noted that a single watermark "cell" commonly encompasses a large number of ink droplets. In high resolution printing, as is commonly used in security documents (e.g. 5000 microdroplets per inch), a single watermark cell may encompass a region of 50 droplets by 50 droplets. In other embodiments, a cell may encompass greater or lesser numbers

of droplets.

[0050] Decoding a watermark requires precise re-registration of the scanned document image, so the watermark cells are located where expected. To facilitate such registration, a calibration signal can be employed.

[0051] An exemplary calibration signal is a geometrical pattern having a known Fourier-Mellin transform. As described in patent 5,862,260, when a known pattern is transformed into the Fourier domain, and then further transformed into the Fourier-Mellin domain, the transformed data indicates the scale and rotation of the pattern. If this pattern is replicated on a security document that is thereafter scanned (as noted, scanning commonly introduces rotation, and sometimes scaling), the F-M transform data indicates the scale and rotation of the scanned data, facilitating virtual re-registration of the security document image for watermark detection.

[0052] As shown in Fig. 7B, an illustrative geometrical calibration pattern is a block, 3.2 cm on a side. The block comprises a 16 x 16 array of substantially identical tiles, each 2 mm on a side. Each tile, in turn, comprises an 8 x 8 array of component cells.

[0053] As described below, the geometrical calibration pattern in the illustrated embodiment is a *visible* design feature on the security document. Accordingly, unlike the watermark data, the calibration pattern does not have to be limited to a small range of digital numbers in order to keep it substantially hidden among other features of the document. Also unlike the watermark data, the illustrated calibration pattern is not locally scaled in accordance with data hiding attributes of the security document image.

[0054] It is possible to print rectangular grids of grey-scaled ink on a document to serve as a calibration pattern. However, aesthetic considerations usually discourage doing so. Preferable is to realize the calibration pattern in a more traditional art form, such as a seemingly random series of intertwining lines, forming a weave-like pattern that is printed across part or all of the document.

[0055] To create this weave-like calibration pattern, a designer first defines an 8 x 8 cell reference calibration tile. Each cell in the tile is assigned a grey-scale value. In the illustrated embodiment, values within 2-10 percent of each other are used, although this is not essential. An exemplary reference calibration tile is shown in Fig. 8 (assuming 8-bit quantization).

[0056] The Fourier-Mellin transform of a block derived from this reference calibration tile will serve as the key by which the scale and rotation of a scanned security document image are determined.

[0057] There is some optimization that may be done in selecting/designing the pattern of grey-scale values that define the reference calibration tile. The pattern should have a F-M transform that is readily distinguished from those of other design and watermark elements on the security document. One design procedure effects a trial F-M transform of the rest of the security

document design, and works backwards from this data to select a reference calibration tile that is readily distinguishable.

[0058] Once a reference tile pattern is selected, the next steps iteratively define a tile having a weave-like pattern whose local luminance values approximately match the reference tile's grey-scale pattern.

[0059] Referring to Fig. 9A, the first such step is to select points on the bottom and left side edges of the tile where lines are to cross the tile boundaries. The angles at which the lines cross these boundaries are also selected. (In the illustrated embodiment, these points and angles are selected arbitrarily, although in other embodiments, the choices can be made in conformance with an optimizing design procedure.)

[0060] The selected points and angles are then replicated on the corresponding right and top edges of the tile. By this arrangement, lines exiting the top of one tile seamlessly enter the bottom of the adjoining tile at the same angle. Likewise, lines exiting either side of a tile seamlessly join with lines in the laterally adjoining blocks.

[0061] The designer next establishes trial line paths snaking through the tile (Figs. 9B, 9C), linking arbitrarily matched pairs of points on the tile's edges. (These snaking paths are sometimes termed "worms.") Desirably, these paths pass through each of the 64 component cells forming the tile, with the total path length through each cell being within +1/- 30% of the average path length through all cells. (This trial routing can be performed with pencil and paper, but more commonly is done on a computer graphics station, with a mouse, light pen, or other input device being manipulated by the designer to establish the routing.) In the illustrated embodiment, the lines have a width of about 30- 100 microns, and an average spacing between lines of about 100 - 400 microns, although these parameters are not critical.

[0062] Turning next to Fig. 10, the trial tile is assembled with like tiles to form a 16 x 16 trial block (3.2 cm on a side), with a repetitive weave pattern formed by replication of the line pattern defined on the 8 x 8 cell trial tile. This trial block is then converted into grey-scale values. The conversion can be done by scanning a printed representation of the trial block, or by computer analysis of the line lengths and positions. The output is a 128 x 128 array of grey-scale values, each value corresponding to the luminance of a 250 micron cell within the trial block.

[0063] This grey-scale data is compared with grey-scale data provided by assembling 256 of the reference calibration tiles (each an 8 x 8 array of cells) into a 16 x 16 calibration pattern block. In particular, the grey-scale array resulting from the trial block is subtracted from the grey-scale array resulting from the reference block, generating a 128 x 128 array of error values. This error data is used to tweak the arrangement of lines in the trial block.

[0064] In cells of the trial calibration block where the error value is positive, the line is too long. That is, the pattern is too dark in those cells (i.e. it has a low luminance grey-scale value), due to a surplus of line length (i.e. too much ink). By shortening the line length in those cells, their luminance is increased (i.e. the cell is lightened). Shortening can be effected by straightening curved arcs, or by relocating a line's entrance and exit points in a cell so less distance is traversed through the cell.

[0065] Conversely, in cells where the error value is negative, the line is too short. By increasing the line length in such cells, their luminance is decreased (i.e. the cell is darkened). Increasing the line length through a cell can be accomplished by increasing the curvature of the line in the cell, or by relocating a line's entrance and exit points along the boundary of the cell, so more distance is traversed through the cell.

[0066] A computer program is desirably employed to effect the foregoing changes in line routing to achieve the desired darkening or lightening of each cell.

[0067] After line positions in the trial calibration block have been tweaked in this fashion, the trial block is again converted to grey-scale values, and again subtracted from the reference block. Again, an array of error values is produced. The positions of the lines are then further tweaked in accordance with the error values.

[0068] The foregoing steps of tweaking line routes in accordance with error signals, converting anew into grey-scale, and computing new error values, is repeated until the luminance of the resulting weave pattern in the trial block is arbitrarily close to the luminance of the reference block. Four or five iterations of this procedure commonly suffice to converge on a final calibration block.

[0069] (It will be noted that the initial tile pattern created by the designer is done at the tile level — 8 x 8 cells. After the initial trial tile is created, subsequent processing proceeds at the block level (128 x 128 cells). A common result of the iterative design procedure is that the component tiles lose their uniformity. That is, the pattern of lines in a tile at a corner of the final calibration block will generally be slightly different than the pattern of lines in a tile near the center of the block.)

[0070] After the final calibration block pattern has been established as above, the blocks are tiled repetitively over some or all of the security document, and can serve either as a background design element, or as a more apparent element of the design. By printing this weave pattern in an ink color close to the paper substrate color, the patterning is highly unobtrusive. (If a highly contrasting ink color is used, and if the pattern extends over most or all of the security document, it may be desirable to employ a brighter luminance paper than otherwise, since the weave pattern effectively darkens the substrate.)

[0071] As noted in my application 08/649,419, the Fourier-Mellin transform has the property that the same

output pattern is produced, regardless of rotation or scaling of the input image. The invariant output pattern is shifted in one dimension proportional to image rotation, and shifted in another dimension proportional to image scaling. When an image whose F-M transform is known, is thereafter rotated and/or scaled, the degree of rotation and scaling can be determined by observing the degree of shift of the transformed F-M pattern in the two dimensions. Once the rotation and scale are known, reciprocal processing of the image can be performed to restore the image to its original orientation and scale.

[0072] In the above-described embodiment, the calibration block pattern has a known F-M transform. When a security document incorporating such a pattern is scanned (e.g. by a photocopier, a flatbed scanner, a facsimile machine, etc.), the resulting data can be F-M transformed. The known F-M pattern is then identified in the transformed data, and its two-dimensional shift indicates the scale and rotation corruption of the scanned security document data. With these parameters known, misregistration of the security document — including scale and rotation corruption — can be backed-off, and the security document data restored to proper alignment and scale. In this re-registered state, the watermark can be detected. (In alternative embodiments, the original scan data is not processed to remove the scale/rotation effects. Instead, subsequent processing proceeds with the data in its corrupted state, and takes into account the specific corruption factor(s) to nonetheless yield accurate decoding, etc.)

[0073] The just-described calibration pattern and design procedure, of course, are just exemplary, and are subject to numerous modifications. The dimensions can be varied at will. It is not essential that the cell size of the calibration tiles match that of the watermark. Nor do the cells sizes need to be integrally related to each other. Nor does the calibration pattern need to be implemented as lines; other ink patterns can alternatively be used to approximate the grey-scale reference pattern.

[0074] There is no requirement that the lines snake continuously through the tiles. A line can connect to just a single edge point of a tile, resulting in a line that crosses that tile boundary, but no other. Or a line can both begin and end in a single tile, and not connect to any other.

[0075] While darker lines on a lighter background are illustrated, lighter lines on a darker background can alternatively be employed.

[0076] The iterative design procedure can employ the F-M transform (or other transform). For example, the trial block pattern can be transformed to the F-M domain, and there compared with the F-M transform of the reference block. An F-M domain error signal can thus be obtained, and the routing of the lines can be changed in accordance therewith.

[0077] Although the illustrated embodiment tweaked the cell-based grey-scales of the calibration block by changing line curvature and position, other luminance

changing techniques can be employed. For example, the width of the weave lines can be locally changed, or small ink dots can be introduced into certain cell areas.

[0078] The foregoing (and following) discussions contemplate that the watermark and/or calibration pattern is printed at the same time as (indeed, sometimes as part of) the line art on the security document. In many applications it is desirable to provide the calibration pattern on the security document substrate prior to printing. The markings can be ink applied by the manufacturer, or can be embossings applied. e.g., by rollers in the paper-making process. (Such textural marking is discussed further below.) Or, the markings can be applied by the security document printer, as a preliminary printing operation, such as by offset printing. By using an ink color/density that is already closely matched to the underlying tint of the paper stock, the manufacturer of the paper can introduce less tinting during its manufacture. Such tinting will effectively be replaced by the preliminary printing of the watermark/calibration pattern on the blank paper.

[0079] Calibration signals entirely different than those detailed above can also be used. Calibration signals that are optimized to detect rotation, but not scaling, can be employed when scaling is not a serious concern. DCT and Fourier transforms provide data that is readily analyzed to determine rotation. A calibration signal can be tailored to stand out in a typically low-energy portion of the transformed spectrum (e.g. a series of fine lines at an inclined angle transforms to a usually vacant region in DCT space), and the scanned image can be transformed to the DCT/Fourier domains to examine any shift in the calibration signal (e.g. a shift in the spatial frequency representation of the inclined lines).

[0080] In some security documents, the just-described calibration weave is printed independently of the watermark encoding. In other embodiments, the weave serves as the lines whose widths, locations, etc., are modulated by the watermark data, as detailed herein and in application 09/074,034.

[0081] In an illustrative embodiment, the printing of the security document is achieved by intaglio printing. Intaglio is a well known printing process employing a metal plate into which the security document pattern is etched or engraved. Ink is applied to the plate, filling the etched recesses/grooves. Paper is then pressed into the plate at a very high pressure (e.g. 10 - 20 tons), both raised-inking and slightly deforming (texturing) the paper.

[0082] Although ink is commonly used in the intaglio process, it need not be in certain embodiments of the present invention. Instead, the paper texturing provided by the intaglio pressing - alone - can suffice to convey watermark data. (Texturing of a medium to convey watermark information is disclosed in various of my prior applications, including allowed application 08/438,159.)

[0083] To illustrate, an intaglio plate was engraved

(using a numerically controlled engraving apparatus), to a depth of slightly less than 1 mm, in accordance with a 3.2 x 3.2 cm. noise-like block of watermark data. The watermark data was generated as described above (e.g. 128 bits of data, randomly distributed in a 128 x 128 cell array), and summed with a correspondingly-sized block of calibration data (implemented as discrete grey-scaled cells, rather than the line/weave pattern detailed above). In this embodiment, the data was not kept within a small range of digital numbers, but instead was raised to a full 8-bit dynamic range.) Banknote paper was intaglio-pressed into this plate - without ink - yielding a generally flat substrate with a 3.2 x 3.2 cm textured region therein. Only on fairly close inspection was the texturing visible; on casual inspection the paper surface appeared uniform.

[0084] This textured paper was placed - textured extrema down - on the platen of an conventional flatbed scanner (of the sort commonly sold as an accessory for personal computers), and scanned. The resulting image data was input to Adobe's Photoshop image processing software, version 4.0, which includes Digimarc watermark reader software. The software readily detected the watermark from the textured paper, even when the paper was skewed on the scanner platen.

[0085] The optical detection process by which a seemingly blank piece of paper can reliably convey 128 bits of data through an inexpensive scanner has not been analyzed in detail; the degree of localized reflection from the paper may be a function of whether the illuminated region is concave or convex in shape. Regardless of the explanation, it is a remarkable phenomenon to witness.

[0086] A second experiment was conducted with the same engraved plate, this time using transparent ink. The results were similar, although detection of the watermark data was not always as reliable as in the inkless case. The raised transparent ink may serve as light conduit, dispersing the incident illumination in unpredictable ways as contrasted with simple reflection off un-inked paper.

[0087] Experiments have also been conducted using traditional opaque inks. Again, the watermark can reliably be read.

[0088] In addition to the just-described technique for "reading" intaglio markings by a conventional scanner, a variant technique is disclosed in Van Renesse, Optical Inspection Techniques for Security Instrumentation, SPIE Proc. Vol. 2659, pp. 159-167 (1996), and can alternatively be used in embodiments according to the present invention.

[0089] Although intaglio is a preferred technique for printing security documents, it is not the only such technique. Other familiar techniques by which watermarks and calibration patterns can be printed include offset litho and letterpress, as well as inkjet printing, xerographic printing, etc. And, as noted, textured watermarking can be effected as part of the paper-making

process, e.g. by high pressure textured rollers.

[0090] In still other embodiments, the watermark and/or calibration ("information") patterns are not printed on the security document substrate, but rather are formed on or in an auxiliary layer that is laminated with a base substrate. If a generally clear laminate is used, the information patterns can be realized with opaque inks, supplementing the design on the underlying substrate. Or the added information can be encoded in textural form. Combinations of the foregoing can similarly be used.

[0091] To retrofit existing security document designs with information patterns, the existing artwork must be modified to effect the necessary additions and/or tweaks to localized security document luminance and/or texture.

[0092] When designing new security documents, it would be advantageous to facilitate integration of information patterns into the basic design. One such arrangement is detailed in the following discussion.

[0093] Many security documents are still designed largely by hand. A designer works at a drafting table or computer workstation, and spends many hours laying-out minute (e.g. 5 mm x 5 mm) excerpts of the design. To aid integration of watermark and/or calibration pattern data in this process, an accessory layout grid can be provided, identifying the watermark "bias" (e.g. -3 to +3) that is to be included in each 250 micron cell of the security document. If the accessory grid indicates that the luminance should be slightly increased in a cell (e.g. 1%), the designer can take this bias in mind when defining the composition of the cell and include a touch less ink than might otherwise be included. Similarly, if the accessory grid indicates that the luminance should be somewhat strongly increased in a cell (e.g. 5%), the designer can again bear this in mind and try to include more ink than might otherwise be included. Due to the substantial redundancy of most watermark encoding techniques, strict compliance by the designer to these guidelines is not required. Even loose compliance can result in artwork that requires little, if any, further modification to reliably convey watermark and/or calibration information.

[0094] Such "designing-in" of embedded information in security documents is facilitated by the number of arbitrary design choices made by security document designers. A few examples from U.S. banknotes include the curls in the presidents' hair, the drape of clothing, the clouds in the skies, the shrubbery in the landscaping, the bricks in the pyramid, the fill patterns in the lettering, and the great number of arbitrary guilloche patterns and other fanciful designs, etc. All include curves, folds, wrinkles, shadow effects, etc., about which the designer has wide discretion in selecting local luminance, etc. Instead of making such choices arbitrarily, the designer can make these choices deliberately so as to serve an informational — as well as an aesthetic — function.

[0095] To further aid the security document designer, data defining several different information-carrying patterns (both watermark and/or calibration pattern) can be stored on mass storage of a computer a workstation and serve as a library of design elements for future designs. The same user-interface techniques that are employed to pick colors in image-editing software (e.g. Adobe Photoshop) and fill textures in presentation programs (e.g. Microsoft PowerPoint) can similarly be used to present a palette of information patterns to a security document designer. Clicking on a visual representation of the desired pattern makes the pattern available for inclusion in a security document being designed (e.g. filling a desired area).

[0096] In the embodiment earlier-described, the calibration pattern is printed as a visible artistic element of the security document. However, the same calibration effect can be provided subliminally if desired. That is, instead of generating artwork mimicking the grey-scale pattern of the reference calibration block, the reference calibration block can itself be encoded into the security document as small changes in local luminance. In many such embodiments, the bias to localized document luminance due to the calibration pattern is simply added to the bias due to the watermark data, and encoded like the watermark data (e.g. as localized changes to the width or position of component line-art lines, as inserted ink droplets, etc.).

[0097] The uses to which the 128 bits of watermark data can be put in security documents are myriad. Many are detailed in the materials cited above. Examples include postal stamps encoded with their value, or with the zip code of the destination to which they are addressed (or from which they were sent); banknotes encoded with their denomination, and their date and place of issuance; identification documents encoded with authentication information by which a person's identity can be verified; etc., etc.

[0098] The encoded data can be in a raw form - available to any reader having the requisite key data (in watermarking techniques where a key data is used), or can be encrypted, such as with public key encryption techniques, etc. The encoded data can embody information directly, or can be a pointer or an index to a further collection of data in which the ultimate information desired is stored.

[0099] For example, watermark data in a passport need not encode a complete dossier of information on the passport owner. Instead, the encoded data can include key data (e.g. a social security number) identifying a particular record in a remote database in which biographical data pertaining to the passport owner is stored. A passport processing station employing such an arrangement is shown in Fig. 11.

[0100] To decode watermark data, the security document must be converted into electronic image data for analysis. This conversion is typically performed by a scanner.

[0101] Scanners are well known, so a detailed description is not provided here. Suffice it to say that scanners conventionally employ a line of closely spaced photodetector cells that produce signals related to the amount of the light reflected from successive swaths of the document. Most inexpensive consumer scanners have a resolution of 300 dots per inch (dpi), or a center to center spacing of component photodetectors of about 84 microns. Higher quality scanners of the sort found in most professional imaging equipment and photocopiers have resolutions of 600 dpi (42 microns), 1200 dpi (21 microns), or better.

[0102] Taking the example of 300 dpi scanner (84 micron photodetector spacing), each 250 micron region 12 on the security document will correspond to about a 3 x 3 array of photodetector samples. Naturally, only in rare instances will a given region be physically registered with the scanner so that nine photodetector samples capture the luminance in that region, and nothing else. More commonly, the image is rotated with respect to the scanner photodetectors, or is longitudinally misaligned (i.e. some photodetectors image sub-parts of two adjoining regions). However, since the scanner oversamples the regions, the luminance of each region can unambiguously be determined.

[0103] In one embodiment, the scanned data from the document is collected in a two dimensional array of data and processed to detect the embedded calibration information. The scanner data is then processed to effect a virtual re-registration of the document image. A software program next analyzes the statistics of the re-registered data (using the techniques disclosed in my prior writings) to extract the bits of the embedded data.

[0104] (Again, the reference to my earlier watermark decoding techniques is exemplary only. Once scanning begins and the data is available in sampled form, it is straightforward to apply any other watermark decoding technique to extract a correspondingly-encoded watermark. Some of these other techniques employ domain transformations (e.g. to wavelet, DCT, or Fourier domains, as part of the decoding process).)

[0105] In a variant embodiment, the scanned data is not assembled in a complete array prior to processing. Instead, it is processed in real-time, as it is generated, in order to detect embedded watermark data without delay. (Depending on the parameters of the scanner, it may be necessary to scan a half-inch or so of the document before the statistics of the resulting data unambiguously indicate the presence of a watermark.)

[0106] In other embodiments, hardware devices are provided with the capability to recognize embedded watermark data in any document images they process, and to respond accordingly.

[0107] One example is a color photocopier. Such devices employ a color scanner to generate sampled (pixel) data corresponding to an input media (e.g. a dollar bill). If watermark data associated with a security document is detected, the photocopier can take one or

more steps.

[0108] One option is simply to interrupt copying, and display a message reminding the operator that it is illegal to reproduce currency.

[0109] Another option is to dial a remote service and report the attempted banknote reproduction. Photocopiers with dial-out capabilities are known in the art (e.g. patent 5,305,199) and are readily adapted to this purpose. The remote service can be an independent service, or can be a government agency.

[0110] Yet another option is to permit the copying, but to insert forensic tracer data in the resultant copy. This tracer data can take various forms. Steganographically encoded binary data is one example. An example is shown in patent 5,568,268. The tracer data can memorialize the serial number of the machine that made the copy and/or the date and time the copy was made. To address privacy concerns, such tracer data is not normally inserted in all photocopied output, but is inserted only when the subject being photocopied is detected as being a security document. (An example of such an arrangement is shown in Fig. 12.)

[0111] Desirably, the scan data is analyzed on a line-by-line basis in order to identify illicit photocopying with a minimum of delay. If a security document is scanned, one or more lines of scanner output data may be provided to the photocopier's reprographic unit before the recognition decision has been made. In this case the photocopy will have two regions: a first region that is not tracer-marked, and a second, subsequent region in which the tracer data has been inserted.

[0112] Photocopiers with other means to detect not-to-be-copied documents are known in the art, and employ various response strategies. Examples are detailed in U.S. Patents 5,583,614, 4,723,149, 5,633,952, 5,640,467, and 5,424,807.

[0113] Another hardware device that can employ the foregoing principles is a standalone scanner. A programmed processor (or dedicated hardware) inside the scanner analyzes the data being generated by the device, and responds accordingly.

[0114] Yet another hardware device that can employ the foregoing principles is a printer. A processor inside the device analyzes graphical image data to be printed, looking for watermarks associated with security documents.

[0115] For both the scanner and printer devices, response strategies can include disabling operation, or inserting tracer information. (Such devices typically do not have dial-out capabilities.)

[0116] Again, it is desirable to process the scanner or printer data as it becomes available, so as to detect any security document processing with a minimum of delay. Again, there will be some lag time before a detection decision is made. Accordingly, the scanner or printer output will be comprised of two pans, one without the tracer data, and another with the tracer data.

[0117] Many security documents already include visi-

ble structures that can be used as aids in banknote detection (e.g. the seal of the issuer and various geometrical markings). In accordance with a further aspect of the present invention, a security document is analyzed by an integrated system that considers both the visible structures and watermark-embedded data.

[0118] Visible security document structures can be sensed using known pattern recognition techniques. Examples of such techniques are disclosed in U.S. Patents 5,321,773, 5,390,259, 5,533,144, 5,539,841, 5,583,614, 5,633,952, 4,723,149, 5,692,073, and 5,424,807 and laid-open foreign applications EP 649,114 and EP 766,449.

[0119] In photocopiers (and the like) equipped to detect both visible structures and watermarks from security documents, the detection of either can cause one or more of the above-noted responses to be initiated (Fig. 12).

[0120] Again, scanners and printers can be equipped with a similar capability - analyzing the data for either of these security document hallmarks. If either is detected, the software (or hardware) responds accordingly.

[0121] Identification of security documents by watermark data provides an important advantage over recognition by visible structures - it cannot so easily be defeated. A security document can be doctored (e.g. by white-out, scissors, or less crude techniques) to remove/obliterate the visible structures. Such a document can then be freely copied on either a visible structure-sensing photocopier or scanner/printer installation. The removed visible structure can then be added back in via a second printing/photocopying operation. If the printer is not equipped with security document-disabling capabilities, image-editing tools can be used to insert visible structures back into image data sets scanned from such doctored documents, and the complete document can then be freely printed. By additionally including embedded watermark data in the security document, and sensing same, such ruses will not succeed.

[0122] (A similar ruse is to scan a security document image on a non-security document-sensing scanner. The resulting image set can then be edited by conventional image editing tools to remove/obliterate the visible structures. Such a data set can then be printed - even on a printer/photocopier that examines such data for the presence of visible structures. Again, the missing visible structures can be inserted by a subsequent printing/photocopying operation.)

[0123] Desirably, the visible structure detector and the watermark detector are integrated together as a single hardware and/or software tool. This arrangement provides various economies, e.g., in interfacing with the scanner, manipulating pixel data sets for pattern recognition and watermark extraction, electronically re-registering the image to facilitate pattern recognition/watermark extraction, issuing control signals (e.g. disabling) signals to the photocopier/scanner,

etc.

[0124] While the foregoing apparatuses are particularly concerned with counterfeit deterrence, the embedded markings can also serve other functions. Examples include banknote processing machines that perform denomination sorting, counterfeit detection, and circulation analysis functions. (I.e., banknotes with certain markings may be distributed through known sources, and their circulation/distribution can subsequently be monitored to assist in macro-economic analyses.)

[0125] From the foregoing, it will be recognized that various embodiments according to the present invention provide techniques for embedding multi-bit binary data in security documents, and provide for the reliable extraction of such data even in the presence of various forms of corruption (e.g. scale and rotation).

[0126] Having described and illustrated the principles of my invention with reference to several illustrative embodiments, it will be recognized that these embodiments are exemplary only and should not be taken as limiting the scope of my invention. Guided by the foregoing teachings, it should be apparent that other watermarking, decoding, and anti-counterfeiting technologies can be substituted for, and/or combined with, the elements detailed above to yield advantageous effects. Other features disclosed in my earlier applications can similarly be employed in embodiments of the technology detailed herein. (Thus, I have not here belabored application of each of the techniques disclosed in my earlier applications - e.g. use of neural networks for watermark detectors - to the present subject matter since same is fairly taught by reading the present disclosure in the context of my earlier work.)

[0127] While the technology has been described with reference to embodiments employing regular rectangular arrays of cells, those skilled in the art will recognize that other arrays - neither rectangular nor regular - can alternatively be used.

[0128] While the embodiments have described the calibration patterns as adjuncts to digital watermarks - facilitating their detection, such patterns have utility apart from digital watermarks. One example is in re-registering scanned security document image data to facilitate detection of visible structures (e.g. detection of the printed seal of the issuer, using known pattern recognition techniques). Indeed, the use of such calibration patterns to register both watermark and visible structure image data for recognition is an important economy that can be gained by integration a visible structure detector and a watermark detector into a single system.

[0129] Although security documents have most commonly been printed (e.g. cotton/linen), other substrates are gaining in popularity (e.g. synthetics, such as polymers) and are well (or better) suited for use with the above-described techniques.

[0130] The embodiments detailed above can be implemented in dedicated hardware (e.g. ASICs), programmable hardware, and/or software.

[0131] Several patents and patent publications were noted above. The artisan is presumed to be familiar with the technologies described therein, and to be capable of integrating same with the elements particularly detailed in this disclosure. To the extent possible, the disclosures of those documents is incorporated herein by reference.

[0132] In view of the many possible embodiments to which the principles of the above-described technology may be put, it should be recognized that the detailed embodiments are illustrative only and should not be taken as limiting the scope of my invention. Rather, I claim as my invention all such embodiments as may come within the scope and spirit of the following claims and equivalents thereto.

Claims

1. A security document including at least a substrate, a digital watermark, and a calibration pattern, the calibration pattern having a known transform facilitating decoding of the digital watermark from scan data corresponding to said document. 20
2. The document of claim 1 in which the calibration pattern has a known Fourier-Mellin transform facilitating decoding of the digital watermark from scan data corresponding to said document. 25
3. A negotiable financial instrument in accordance with claim 1. 30
4. A passport in accordance with claim 1.
5. The document of claim 1 in which the calibration pattern has a visible weave-like appearance. 35
6. The document of claim 1 in which the calibration pattern extends across most of the security document. 40
7. The document of claim 1 in which at least one of the digital watermark and calibration pattern are realized by texturing of the document, and not printing. 45
8. The document of claim 1 in which both the digital watermark and the calibration pattern are realized by texturing of the document, and not printing. 50
9. The document of claim 1 in which each one square millimeter zone of the calibration pattern has an 8-bit grey-scale value within twenty percent of the value of all adjoining one square millimeter zones of the calibration pattern. 55
10. The document of claim 1 in which each 250 micron by 250 micron zone of the calibration pattern has an 8-bit grey-scale value within twenty percent of the value of all adjoining 250 micron by 250 micron

zones of the calibration pattern.

11. The document of claim 1 in which the calibration pattern comprises a tiled pattern, each tile including a line following an apparently random route.
12. A method of producing a security document, comprising:
 - providing multi-bit binary data;
 - forming a pattern of recesses in a metal member, said pattern having the multi-bit binary data encoded therein; and
 - pressing a substrate into said metal member to texture the substrate in accordance with said binary data.
13. The method of claim 12 further including inking the recesses prior to the pressing.
14. The method of claim 12 in which the pattern has the multi-bit binary data steganographically encoded therein, wherein a human viewer of the texture on the substrate is not alerted that it conveys multi-bit binary data.
15. A method of processing a security document, comprising:
 - generating scan data corresponding to said document;
 - performing a domain transformation on the scan data to obtain corruption data related thereto; and
 - detecting embedded information from the processed scan data, said detecting employing said corruption data.
16. The method of claim 15 in which:
 - the domain transformation is a Fourier-Mellin transform;
 - the corruption data is a scale or rotation factor; and
 - the detecting includes compensating the scan data in accordance with the corruption data.
17. The method of claim 15 in which the embedded information comprises multi-bit digital data.
18. The method of claim 15 which further includes using the detected information to access additional information in a database.
19. Apparatus for use with security documents, comprising:
 - a scanner for scanning a security document

and producing scan data corresponding thereto;

a processor responsive to encoding on the security document for determining rotation of the scan data from a reference state and producing output data corresponding thereto; 5

a visible structure detector; and

a steganographic watermark detector;

said output data permitting the detectors to detect visible structures and steganographic 10 watermarks associated with the security document notwithstanding at least certain rotations of the scan data.

15

20

25

30

35

40

45

50

55

FIG. 1A (Prior Art)




FIG. 1B (Prior Art)




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7A




FIG. 7B




FIG. 8

| | | | | | | | |
|-----|-----|-----|-----|-----|-----|-----|-----|
| 204 | 211 | 212 | 214 | 213 | 207 | 215 | 214 |
| 204 | 215 | 202 | 205 | 209 | 205 | 213 | 202 |
| 212 | 207 | 203 | 214 | 203 | 206 | 202 | 215 |
| 209 | 201 | 211 | 201 | 212 | 204 | 200 | 203 |
| 208 | 204 | 212 | 206 | 207 | 203 | 205 | 202 |
| 209 | 214 | 207 | 207 | 211 | 201 | 206 | 213 |
| 208 | 212 | 206 | 211 | 213 | 208 | 206 | 213 |
| 209 | 208 | 202 | 202 | 205 | 205 | 205 | 211 |

REFERENCE GREY-SCALE CALIBRATION TILE

FIG. 12

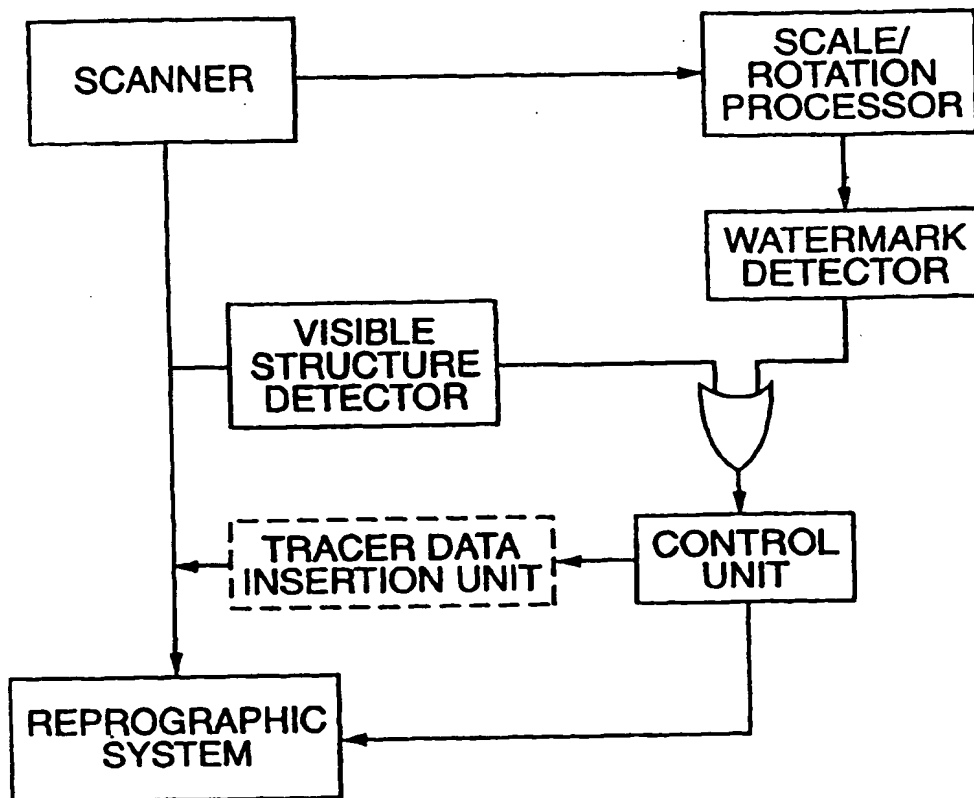


FIG. 9A




FIG. 9B




FIG. 9C

